

CLAIMS

What is claimed is:

1. A method comprising:

 storing in a data structure information representing a plurality of clients on a network; and

 dynamically compressing the data structure based on network proximity information relating to the clients.

2. A method as recited in claim 1, wherein the data structure comprises a network address of each of the clients and the network proximity information.

3. A method as recited in claim 1, wherein said dynamically compressing the data structure based on the network proximity information comprises:

 detecting when proximity measurements for at least two clients which share a network address prefix are within a predetermined range of each other; and

 in response to the proximity measurements being within the predetermined range of each other for the at least two clients, merging entries for the at least two clients in the data structure.

4. A method as recited in claim 1, further compressing losslessly decompressing the data structure.

5. A method as recited in claim 4, wherein said losslessly decompressing the data structure comprises splitting a merged entry in the data structure representing at least two clients into a plurality of separate entries.
6. A method as recited in claim 4, wherein said losslessly decompressing the data structure comprises decompressing the data structure in response to a detected change in network conditions.
7. A method as recited in claim 1, wherein the data structure comprises a network address of each of the clients.
8. A method as recited in claim 7, further comprising:
 - generating an encoded value from each of the network addresses;
 - storing each of the encoded values in the data structure in association with a corresponding network address; and
 - modifying one or more of the encoded values to indicate entries in the data structure that have been merged.
9. A method as recited in claim 1, wherein the data structure comprises the network proximity information.
10. A method as recited in claim 9, wherein the method is performed in a content delivery agent configured to deliver content to the plurality of clients.
11. A method as recited in claim 9, wherein the method is performed in a request routing agent configured to route content requests from any of the plurality of clients.

12. A method as recited in claim 1, further comprising reporting information from the data structure to a request routing server.

13. A method comprising:

receiving information relating to a plurality of clients on a network, the information including a plurality of network addresses and a corresponding plurality of masks, each of the network addresses representing one or more of the clients, each of the masks being indicative of a network proximity-based compression state of received information relating to the clients;

storing the received information relating to the plurality of clients in a data structure; and

dynamically compressing the data structure based on the masks.

14. A method as recited in claim 13, wherein said method is performed in a request routing server.

15. A method as recited in claim 14, wherein the information representing a plurality of clients is received from a plurality of content delivery agents on the network.

16. A method as recited in claim 13, wherein said receiving information comprises receiving a plurality of masks for each of the network addresses, each of the masks being indicative of a proximity-based compression state for a corresponding network address, each of the masks corresponding to a different one of a plurality of request routing agents.

17. A method as recited in claim 13, further comprising generating a master mask for each of the network addresses, based on the plurality of masks;

wherein said dynamically compressing the data structure based on the masks comprises dynamically compressing the data structure based on the master masks.

18. A method as recited in claim 17, wherein said dynamically compressing the data structure comprises, if the master mask is identical for two or more entries in the data structure, then merging the two or more entries in the data structure.

19. A method as recited in claim 13, further comprising losslessly decompressing the data structure.

20. A method as recited in claim 19, wherein said losslessly decompressing the data structure comprises splitting a merged entry in the data structure representing at least two of the clients into a plurality of separate entries.

21. A method as recited in claim 19, wherein said losslessly decompressing the data structure comprises decompressing the data structure in response to a detected change in network conditions.

22. A method of storing client information in a content delivery network, the method comprising:

creating a first data structure in a content delivery agent associated with a plurality of clients, so that the data structure includes a separate entry for each of the clients;

storing in the entry for each client a network address of the client, each network address including a network address prefix;

generating a bitmap mask for each client based on the network address of each client;

storing the bitmap mask for each client in the corresponding entry in the first data structure;

repeatedly measuring a proximity between the agent and each of the plurality of clients;

storing the measured proximity for each client in the corresponding entry in the first data structure;

detecting when the measured proximities are within a predetermined range of each other for at least two clients which have the same network address prefix;

in response to the measured proximities being within the predetermined range of each other for the at least two clients, merging the entries for the at least two clients in the first data structure, including

using the network address of one of the entries being merged to represent both of the entries being merged, and

generating a single bitmap mask to represent the entries being merged as a logic OR of the individual bitmap masks of the entries being merged; and

for each of the entries in the first data structure, reporting the corresponding network address and bitmap mask to a request routing server.

23. A method as recited in claim 22, wherein the network address is an IP address.

24. A method as recited in claim 22, wherein the bitmap mask for each client is an encoded representation of a predetermined number of lowest order bits of the network address of the client.
25. A method as recited in claim 24, wherein the predetermined number of lowest order bits of the network address represents a single host.
26. A method as recited in claim 22, further comprising:
 - determining whether the measured proximities for at least two of the clients represented by a merged entry in the first data structure fall outside a predetermined range of each other; and
 - in response to the measured proximities for at least two of the clients represented by a merged entry in the first data structure falling outside the predetermined range of each other, splitting the merged entry into a plurality of separate entries in the first data structure.
27. A method as recited in claim 22, further comprising:
 - creating a second data structure in the request routing server, the second data structure including a plurality of entries, each representing one or more of the clients, each entry in the second data structure including a network address, a plurality of bitmap masks, and a master bitmap mask, each of the plurality of bitmap masks having been reported to the request routing server by a separate one of a plurality of content delivery agents, each of the plurality of bitmap masks representing the network address of one or more of the clients.

28. A method as recited in claim 27, wherein for each entry in the second data structure, the master bitmap mask represents a combination of the plurality of bitmap masks for said entry.

29. A method as recited in claim 28, further comprising:

if the master bitmap mask is identical for two or more entries in the second data structure, then merging the two or more entries in the second data structure.

30. A content delivery system comprising:

a plurality of agents to deliver content to a plurality of clients via a network, each agent configured to acquire network proximity information relating to each of the clients, to store in a first data structure information representing the plurality of clients, and to dynamically compress the first data structure based on the network proximity information; and

a server to receive from the agents information relating to the clients, including the network proximity information, and to cause redirection of content requests from the clients to appropriate ones of the agents, based on the network proximity information.

31. A content delivery system as recited in claim 30, wherein the information relating to the clients received by the server comprises a plurality of network addresses and a corresponding plurality of masks from the compressed first data structure, each of the network addresses representing one or more of the clients, each of the masks being indicative of a proximity-based compression state for a corresponding network address.

32. A content delivery system as recited in claim 31, wherein the server is further configured to store the information relating to the clients in a second data structure and to dynamically compress the second data structure based on the masks.

33. A processing system comprising:

 a processor;

 a network communication device to enable the processing system to communicate with a plurality of clients over a network; and

 a storage device containing instructions which, when executed by the processor, cause the processing system to perform a process that includes
 creating a data structure that contains information representing the plurality of clients,

 acquiring network proximity information relating to the clients, and
 dynamically compressing the data structure based on the network proximity information.

34. A processing system as recited in claim 33, wherein the data structure comprises a network address of each of the clients and the network proximity information.

35. A processing system as recited in claim 34, wherein said dynamically compressing the data structure based on the network proximity information comprises:

 detecting when proximity measurements for at least two clients which share a network address prefix are within a predetermined range of each other; and

in response to the proximity measurements being within the predetermined range of each other for the at least two clients, merging entries for the at least two clients in the data structure.

36. A processing system as recited in claim 35, wherein said process further comprises:

forming a mask to represent each of the network addresses;

storing the masks in the data structure; and

modifying one or more of the masks to indicate entries in the data structure that have been compressed.

37. A processing system as recited in claim 33, wherein the processing system comprises a content delivery agent configured to deliver content to any of the plurality of clients, and said process is performed by the content delivery agent.

38. A processing system as recited in claim 33, wherein the processing system comprises a request routing agent configured to route content requests from any of the plurality of clients, and said process is performed by the request routing agent.

39. A processing system as recited in claim 33, wherein said process further comprises reporting information from the data structure to a request routing server.

40. A processing system as recited in claim 33, wherein said process further comprises losslessly decompressing the data structure.

41. A processing system as recited in claim 40, wherein said losslessly decompressing the data structure comprises splitting a merged entry in the data structure representing at least two of the clients into a plurality of separate entries.

42. A processing system as recited in claim 40, wherein said losslessly decompressing the data structure comprises decompressing the data structure in response to a detected change in network conditions.

43. A network caching device comprising:

 a processor;

 a network communication device to enable the network caching device to communicate with a plurality of clients over a network; and

 a storage device containing:

 a first set of instructions which configure the network caching device to cache network-based content and to deliver the cached content to at least one of the clients via the network in response to a request; and

 a second set of instructions which configure the network caching device to perform a process that includes:

 creating a data structure associated with the plurality of clients, so that the data structure includes a separate entry for each of the clients;

 storing in the entry for each client a network address of the client, each network address including a network address prefix;
 repeatedly measuring a network proximity for each of the plurality of clients;

storing the measured network proximity for each of the plurality of clients in the corresponding entry in the data structure;

detecting when the measured network proximities are within a predetermined range of each other for at least two clients which have the same network address prefix; and

in response to the measured network proximities being within the predetermined range of each other for the at least two clients, merging the entries for the at least two clients in the data structure.

44. A network caching device as recited in claim 43, wherein said process further comprises:

generating a mask for each of the plurality of clients based on the network address of said client; and

storing the mask for each of the plurality of clients in the corresponding entry in the data structure, prior to said merging;

wherein said merging comprises:

using the network address of one of the entries being merged to represent both of the entries being merged, and

generating a single mask to represent the entries being merged as a combination of the individual masks of the entries being merged.

45. A network caching device as recited in claim 44, wherein the mask for each of the clients is an encoded representation of a predetermined number of lowest order bits of the network address of the client.

46. A network caching device as recited in claim 44, wherein the predetermined number of lowest order bits of the network address represents a single host.

47. A network caching device as recited in claim 43, wherein the network address of each of the clients is an IP address.

48. A network caching device as recited in claim 43, wherein said process further comprises, for each of the entries in the data structure, reporting the corresponding network address and mask to a request routing server.

49. A network caching device as recited in claim 43, wherein said process further comprises:

determining whether the measured proximities for at least two of the clients represented by a merged entry in the data structure fall outside a predetermined range of each other; and

in response to the measured proximities for at least two of the clients represented by a merged entry in the data structure falling outside the predetermined range of each other, splitting the merged entry into a plurality of separate entries in the data structure.

50. An apparatus comprising:

means for storing in a data structure information representing a plurality of clients on a network;

means for acquiring network proximity information relating to the clients; and

means for dynamically compressing the data structure based on the network proximity information.